REMARKS

Claims 1-44 were pending in the present application. Claims 1, 3-5, 8, 11, 12, 15-18, 21, 24-26, 28, 31, 33, 34, 37, 40, and 44 have been amended. Accordingly, claims 1-44 remain pending in the application.

In paragraph 3 of the present Office action, the Examiner asserts the specification is replete with terms that are not clear, concise, and exact. Applicant respectfully disagrees. More particularly, the Examiner asserts in the sole example, that nowhere in the specification are the specifics of terms such as proxy read to own (PRTO) packets and proxy invalidate (PI) packets described. Applicant respectfully submits the specifics of PRTO and PI packets are described in the specification in a number of places. For example, paragraphs [0189] and [0191], among others.

Applicant notes that due to the complexity of the system, the written description is lengthy and complex to fulfill the enablement and best mode requirements. Accordingly, where appropriate, in the following remarks Applicant has attempted to explain some of the complex concepts used in the specification.

Claims 1-44 are rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the enablement requirement. The Examiner asserts the claim language contradicts itself because the specification defines an "owner" of a coherency unit to be responsible for providing data to another client which requests that coherency unit, and the claim language (e.g., claim 1) recites ignoring the address packet...

Applicant respectfully disagrees. Applicant asserts the definition of the "owner" does not preclude other operation. Moreover, the definition language merely states the owner has responsibility for that function. It does not state the owner provides the data each and every time a request is received. To the contrary, the definition is merely defining who is responsible for providing the data in given situations. However, the

above notwithstanding, Applicant has amended the claims for clarity.

In addition, the Examiner's objected to the statement in paragraph [0189] "While the embodiment illustrated in Fig. 22 uses different types of packets for gM and non-gM nodes, other embodiments may use the same type of packets in all nodes." However, Applicant submits this is just stating that in the embodiment of FIG. 22 a different packet type was used for a given transaction in gM and non-gM nodes, but in other embodiments the same packet type may be used in gM and non-gM nodes. There is nothing indefinite about this. For example, in one embodiment, a PRTOM packet may be used in gM nodes and PRTO packets may be used in non-gM nodes, but in other embodiments only one of those packet types may be used in either gM or non-gM nodes.

The Examiner again asserts the specification does not have any specifics regarding terms such as proxy read to own (PRTO) packets and proxy invalidate (PI) packets. Applicant directs the Examiner to paragraphs [0189], [0191], [0250], [0255], [0260], [0261], and [0272] in which the various limitations recited in the claims are discussed and supported.

Accordingly, from the discussion above, Applicant respectfully requests the Examiner withdraw the rejection under 35 U.S.C. 112, first paragraph.

Claims 1-44 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. Applicant respectfully traverses this rejection.

More particularly, the Examiner asserts the use of a first address packet and a second address packet are not described anywhere. Applicant respectfully disagrees. First, the summary (paragraph [0011]) of Applicant's disclosure uses that language and thus the claim language is broadly supported. In addition, as disclosed in multiple places in Applicant's written description, more than one kind of address packet is sent. For example, if the node is a gM node, one type of address packet is sent (e.g., PRTOM), and

if the node is not gM a second kind of packet is sent (PRTO). Applicant submits it perfectly acceptable to use a "first" and a "second" of anything when two or more are used. Thus, the limitations in the claims are supported and are definite.

The minor typographic errors in the claims (e.g., claims 1, 26 and 28) have been corrected.

In regard to the Examiner's question whether any of the devices may not be active, the answer is unequivocally yes. The plurality of devices may include, as shown in FIG. 1, and FIG. 19 and described at least in paragraph [0052-0053], processing subsystems (active device), memory subsystems (non-active devices), and I/O subsystems (active devices).

Accordingly, Applicant respectfully requests the Examiner withdraw the rejection under 35 U.S.C. 112, second paragraph.

Claims 1-44 are rejected under 35 U.S.C. 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. Specifically, the Examiner contends the claims are indefinite because an active device can't be an owner if the coherency unit is not in the gM state. Although Applicant respectfully traverses this rejection, Applicant has amended the claims for clarity.

However, Applicant notes the definition of the "owner" given in the written description does not preclude other operation. Moreover, the definition language merely states the owner has responsibility for that function. It does not state the owner provides the data each and every time a request is received. To the contrary, the definition is merely defining who is responsible for providing the data in given situations.

Applicant also notes it is possible for an active device to have ownership state for a coherency unit that is in the gS state. This is described at paragraph [0253] lines 11-12

and paragraph [0254] lines 25-30. Thus, the claims are definite, since in the case where the global access state is gM, the active device could be the owner, and if the access state is not gM, the active device is not the owner. This language is not indefinite.

Accordingly, Applicant respectfully requests the Examiner withdraw the rejection under 35 U.S.C. 112, second paragraph.

Claims 1-44 stood rejected under U.S.C. 103(a) as being unpatentable over Hagersten, et al. (U.S. Patent Number 5,940,860) (hereinafter "Hagersten"). Applicant respectfully traverses this rejection.

Applicant's claim 1, as amended, recites a system comprising:

> "a node including one or more active devices, an interface, and an address network configured to transmit address packets between the one or more active devices and the interface; and
> an additional node coupled to the node by an inter-node network, wherein the additional node includes an additional address network;
> wherein a given active device having an ownership responsibility for a coherency unit is configured to respond to certain access right requests;
> wherein in response to receiving from the additional node via the inter-node network, a coherency message requesting an access right to a coherency unit, the interface is configured to send a first type of address packet on the address network if a global access state of the coherency unit *in the node* is a modified state and to send a second type of address packet on the address network if the global access state of the coherency unit *in the node* is *not* the modified state; and
> wherein if the given active device has an ownership responsibility for the coherency unit, the given active device is configured to ignore the second type of address packet and to respond to the first type of address packet." (Emphasis added)

The Examiner asserts Hagersten teaches several of the above limitations based upon the rejections discussed above. However, as Applicant has pointed out, some of the Examiner's rejections under 35 U.S.C. 112 appear to be based upon faulty characterizations of Applicant's disclosure. In addition, Applicant disagrees with the Examiner's characterization of Hagersten. Applicant submits although Hagersten does

disclose a system in which global access states are used, Hagersten does not teach an interface in the node sending different types of packets in the node dependent on whether the requested coherency unit is in the modified state in that node.

To the contrary, as illustrated in FIG. 2 of Hagersten, a computer node (e.g., 100) is connected to other computer nodes (e.g., 150, 160, 170) via links such as 220, 222, etc. A central coherence transformer 200 facilitates communication between the nodes. (See Hagersten; col. 5, line 63-col. 6, line 65). In addition, Hagersten uses a different request protocol since the system components and their arrangement is different. For example, as Hagersten discloses in FIG. 8, there are no proxy packets. Hagersten uses remote requests (e.g., RRTO) to communicate requests form inside the node to the coherence transformer 200. Specifically, at col. 10, lines 10-44, Hagersten discloses

> "A remote RTO (RRTO) memory access request is typically issued by an RTO progenitor after that RTO progenitor finds out, by ascertaining the state of the Mtag received from memory module 110, that the state of the Mtag is insufficient to service the current RTO request. Insufficient Mtag states in this case may be gS or gI, i.e., there may be a shared or exclusive copy of the requested memory block existing externally. If the RRTO is issued by the RTO progenitor responsive to a gM Mtag, coherence transformer understands this to be an error condition (since state gM indicates that the internal domain, not the external domain, currently has the exclusive copy of the requested memory block) and may request the RRTO progenitor to retry to obtain the exclusive copy from the internal domain.
> If the RRTO is issued by the RTO progenitor responsive to a gS Mtag, coherence transformer 200 may respond to this RRTO command by invalidating external shared copy or copies, obtaining the latest copy of the requested memory block either from the external domain or the internal domain, invalidating all internal shared copy or copies, and returning that copy to the RRTO progenitor via the RTOR_data response. If the RRTO is issued by the RTO progenitor responsive to a gI Mtag, coherence transformer 200 may respond to this RRTO command by obtaining the external exclusive copy, invalidating that external exclusive copy, and returning that copy to the RRTO progenitor via the RTOR_data response. Further, coherence transformer 200 may perform a write back to memory module 110 to change the state of the Mtag corresponding to the requested memory block to gM via the RTOR response. If the RRTO request is erroneous, e.g., requesting a non-existent memory block, coherence transformer 200 may reply with a RTOR_nack response, signifying that the RRTO request is not acknowledged and needs to be

retried by the RRTO progenitor." (Emphasis added)

From the above, it is clear that the progenitor (e.g., processor 102) of the request (e.g., RTO) directly issues the request as a remote request (RRTO) to the coherence transformer 200. This is not the same an interface that is separate from the active device issuing one type of proxy packet in response to receiving a request for a coherency unit in a modified state, and issuing another type of proxy packet in response to receiving a request for a coherency unit not in a modified state.

In addition, Applicant submits Hagersten uses the coherence transformer as a centralized unit that is separate from and external to each of the nodes and that receives requests and controls the transactions external to each node. For example, at col. 11, line 53 - col. 12, line 36

> "FIG. 10 illustrates, in one embodiment of the present invention, selected transactions performed by coherence transformer 200 in response to remote memory access requests on common bus 108. Referring now to FIG. 10, when a remote memory access request is issued by one of the internal bus entities on common bus 108, this remote memory access request is forwarded to all bus entities, including coherence transformer 200. The remote request may be, however, ignored by all internal bus entities, e.g., processor 102. Responsive to the remote request, coherence transformer 200 ascertains the current state of the Mtag (included in the remote request) to determine whether one of the external devices has an appropriate copy of the requested memory block for responding to the remote memory access request on common bus 108.
> Remote Request to Own (RRTO)
> If the remote memory access request is a request for an exclusive copy of a memory block (a RRTO) and the current Mtag state is gM, coherence transformer 200 understands this to be an error condition (since state gM indicates that the internal domain, not the external domain, currently has the exclusive copy of the requested memory block) and may request the RRTO progenitor to retry to obtain the exclusive copy from the internal domain.
> If the RRTO is issued by the RTO progenitor responsive to a gS Mtag, coherence transformer 200 may respond to this RRTO command by invalidating external shared copy or copies by issuing the X-protocol invalidate command XINV to request all external devices to invalidate their shared copies. Coherence transformer 200 may either broadcast the X-protocol commands or may simply direct the X-protocol command to the appropriate external device(s) if there is provided logic, either in

hardware or software, with coherence transformer 200 for keeping track of the locations and types of memory blocks cached.
When all external copies have been invalidated (confirmed by the receipt of the X-protocol XINV_nack response) coherence transformer 200 may then obtain the latest copy of the requested memory block from the internal domain and invalidate all internal shared copy or copies. In one embodiment, coherence transformer 200 may obtain the latest copy of the requested memory block from the internal domain and invalidate all internal shared copy or copies by issuing a RTO request to common bus 108. Upon receiving the requested copy from the internal domain (via the RTO_data response), coherence transformer 200 may write back the copy to memory module 110 along with the appropriate Mtag, i.e., gM in this case, via the RTOR response. Thereafter, coherence transformer 200 may provide the requested copy to the RRTO progenitor via the RTOR_ata response."

From the foregoing, the coherence transformer does not issue a different type of address packet based upon whether the coherence unit is in a modified state in a node. Indeed, coherence transformer issues X commands to obtain the requested data from a given node and then returns the data to the requestor. So, although similar in some respects, Hagersten is a different type of system with different request packets and protocols.

Accordingly, Applicant submits Hagersten does not teach or suggest "a node including one or more active devices, an interface," or "wherein in response to receiving from the additional node via the inter-node network, a coherency message requesting an access right to a coherency unit, the interface is configured to send a first type of address packet on the address network if a global access state of the coherency unit in the node is a modified state and to send a second type of address packet on the address network if the global access state of the coherency unit in the node is not the modified state;" or "wherein if the given active device has an ownership responsibility for the coherency unit, the given active device is configured to ignore the second type of address packet and to respond to the first type of address packet" as recited in Applicant's claim1.

Thus, Applicant submits claim 1 along with its dependent claims, patentably distinguishes over Hagersten for the reasons given above.

Applicant's claims 15, 31, and 44 include features that are similar to the features recited in claim 1. Thus, Applicant submits claims 15, 31, and 44 along with their respective dependent claims, patentably distinguish over Hagersten for at least the reasons given above.

CONCLUSION

Applicant submits the application is in condition for allowance, and an early notice to that effect is requested.

If any fees are due, the Commissioner is authorized to charge said fees to Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C. Deposit Account No. 501505/5181-95101/SJC.

Respectfully submitted,

/ Stephen J. Curran /
Stephen J. Curran
Reg. No. 33,929
ATTORNEY FOR APPLICANT(S)

Meyertons, Hood, Kivlin, Kowert, & Goetzel, P.C.
P.O. Box 398
Austin, TX 78767-0398
Phone: (512) 853-8800

Date: July 9, 2007